STOCHASTIC METHODS FOR PORTFOLIO CREDIT DERIVATIVES

Lutz Schloegl

Fixed Income Quantitative Research, Lehman Brothers International (Europe), 25 Bank Street, London E14 5LE, United Kingdom
Email: `luschloe@lehman.com`

Abstract

Credit derivatives are an important meeting ground for actuarial and financial mathematics. This article is a brief introduction to the pricing of portfolio credit derivatives. We survey some of the stochastic methods currently used. These are illustrated with several of the main applications in portfolio credit derivatives such as the pricing of CDO and CDO^2 tranches.

1. CREDIT AT THE INTERSECTION BETWEEN DERIVATIVES AND INSURANCE

Credit derivatives occupy a unique position at the intersection between derivatives and insurance. The most liquid and basic credit derivative, the default swap, is an insurance contract between two counterparties on the credit risk of a reference entity. The *protection buyer* pays a regular premium until default or maturity of the trade, which is known as the *default swap spread* and is quoted on an annualized basis in basis points, i.e. hundredths of a percent of the trade notional. In return, the *protection seller* protects the buyer against the economic loss on the reference entity's bonds in the event of a default. At default, the contract is either subject to cash or physical settlement. In the case of physical settlement, the protection buyer delivers defaulted bonds to the seller and receives their par value in return. In the case of cash settlement, the protection seller pays the difference between par and the bonds' observed recovery rate, i.e. post-default price to the buyer. This is a typical insurance contract. In return for a (relatively) small premium, the protection buyer insures against a rare, but potentially large loss.

The digital nature of credit payouts highlights risks that are not so central to other derivative contracts, this has caused the market to evolve. Maturity mismatches are an important source of risk, and the market has evolved mitigation mechanisms against this. Protection on a new default swap contract commences at $T + 1$ *calendar* days after the trade date T. This is in contrast to other markets where settlement periods are usually expressed in business days to facilitate clearing and other back-office operations. This is particularly important, as credit relevant information is fairly often revealed when markets are closed. Similarly, default swaps are traded to fixed maturity dates. The so-called “IMM” dates are the 20th of March, June, September and December. This reduces maturity mismatches between different long and short positions and significantly facilitates the management of a default swap book.

Default swaps derive their importance not only from their role as insurance contracts, they are also the basic hedging tool for more complex credit derivatives. The exposures stemming from synthetic CDO tranches, CDO^2 trades, default swaptions, etc are all dynamically managed by credit derivative dealers using default swaps. In the context of a trading book, positions are continually marked to market. This means that sensitivities to spread movements are particularly important, more so as a derivative position can mutate from asset to liability and vice versa as the market moves. The nature of counterparty risk also becomes very different: contracts do not just cancel if the premium is no longer paid. Rather, the contract is marked to market and needs to be unwound. If the non-defaulting party is in the money, the market value of the contract becomes an unsecured claim on the defaulting counterparty. In particular, the seller of protection is also exposed to counterparty risk if the market tightens.

The standard approach to bootstrapping credit curves is to take a completely reduced-form view of the default event. The default time τ is a random variable with a distribution modelled via the hazard rate λ by specifying the conditional default probability as

$$P\left[\tau \le t + \Delta t \middle|\, \tau > t\right] = \lambda(t)\Delta t. \tag{1}$$

The hazard rate λ is closely related to the credit spread. In the simple approach, it is treated as a deterministic function. More sophisticated models specify λ as a stochastic process. The unconditional survival probability $Q(0,t)$ to time T is obtained by integrating equation (1):

$$Q(0,T) = P[\tau > T] = \exp\left(-\int_0^T \lambda(s)ds\right).$$

On default, we assume that the loss is $1 - R$, where R is a fixed recovery of par. The hazard rate is calibrated to default swap spreads, which are the actual market observables. Because spreads aggregate loss likelihood and severity, we need to disentangle these two aspects. The protection leg is priced by integrating

$$\text{Prot} = (1 - R)\int_0^T B(0,t)Q(0,t)\lambda(t)dt. \tag{2}$$

Ignoring the issue of coupon accrual, the premium leg is the price of a risky annuity, also known as the risky PV01 (present value of a basis point). It is obtained by summing over the risky discount factors for the coupon dates $T_1, \ldots, T_n$. Denoting the accrual factor for each coupon period by α_i, we have

$$\text{Prem} = \sum_{i=1}^{n} \alpha_i B(0,T_i)Q(0,T_i). \tag{3}$$

The protection and premium leg values are not uniquely determined by the breakeven market spread because

$$s = \frac{\text{Prot}}{\text{Prem}} \approx \frac{(1-R)\sum_{m=1}^{M} B(0,t_m)\left(Q(0,t_{m-1}) - Q(0,t_m)\right)}{\sum_{i=1}^{n} \alpha_i B(0,T_i) Q(0,T_i)}.$$

In particular, different hazard rate curves can give the same default swap spread. Default probabilities are model-dependent quantities that depend quite strongly on the recovery rate assumptions that are used, and to a lesser extent on the interpolation methodology. Nevertheless, no-arbitrage links spreads, recovery rates, and default probabilities. A very useful approximation, the so-called credit triangle, can be derived by assuming a constant hazard rate and a continuously paid spread:

$$s \approx \lambda(1-R).$$

For example, a spread of 90bp and a recovery rate of 40% imply an annual default probability of approximately $\lambda = 1.5\%$. Using the credit triangle, one can compute the market value of protection bought at a spread of s_0:

$$\text{MTM} = (s - s_0)\sum_{i=1}^{n} \alpha_i B(0,T_i) e^{-\frac{sT_i}{1-R}}.$$

The recovery sensitivity of this MTM is quite low, particularly if the current market default swap spread s has not moved far away from s_0. This is good news if one is worried about marking a default swap book correctly. On the other hand, it implies that default swaps do not actually help us in disentangling default and recovery rate risk. The most certain thing about recovery rates is their uncertainty. The best data sources are the rating agencies or perhaps internal ratings, from a credit derivatives perspective one usually has to make fairly broad assumptions, for example a recovery rate of 40% for senior unsecured debt of investment grade companies.

After default swaps, synthetic CDO tranches are the most broadly traded credit derivatives. The insurance character is similar, the protection seller takes exposure to a band of losses to a given reference portfolio. The band is defined by a lower and an upper strike, K_1 and K_2, which are expressed as a percentage of the total portfolio notional. If L_T is the cumulative percentage loss to the portfolio, the percentage loss to the tranche is

$$L_T^{tr} = \frac{[L_T - K_1]^+ - [L_T - K_2]^+}{K_2 - K_1}. \tag{4}$$

The protection seller receives a spread s on the outstanding notional of the tranche. Once the portfolio losses exceed K_1, the seller makes a protection payment every time there is a loss, the notional of the tranche is reduced, and the tranche spread is only paid out on this reduced notional going forward. An important concept for CDO tranche pricing is the so-called tranche "survival probability" $Q^{tr}(0,T)$, which is the expected outstanding notional of the tranche

$$Q^{tr}(0,T) = 1 - \frac{E\left[[L_T - K_1]^+\right] - E\left[[L_T - K_2]^+\right]}{K_2 - K_1}. \tag{5}$$

With this, the protection and premium legs of the tranche swap become analogous to a single-name default swap, see equations (2) and (3):

$$\text{Prot} = -\int_0^T B(0,t)dQ(0,t)$$
$$\text{Prem} = s\sum_{i=1}^{n} \alpha_i B(0,T_i)Q^{tr}(0,T_i).$$

Equation (5) shows that tranche survival probabilities are effectively call spreads on the cumulative portfolio loss. This non-linearity explains why CDO tranches are correlation instruments, they depend not only on the overall risk in the portfolio (determined by $E\left[L_T\right]$), but also critically on the tendency of different credits to default (and survive) together. Hence, the main effort when developing portfolio credit models for tranche pricing is directed towards modelling the dependence structure between the different reference credits.

2. STOCHASTIC MODELLING TECHNIQUES

Equation (5) also implies that we need to concentrate on modelling cumulative portfolio loss distributions. We fix a time horizon T, each credit j defaults with probability p_j. One modelling framework is the so-called latent variable approach. With each credit j, we associate a random variable A_j, such that the credit defaults if A_j falls below a threshold K_j. The value of K_j is calibrated to the marginal default probability p_j:

$$P\left[A_j \leq K_j\right] = p_j. \tag{6}$$

The marginal distribution of A_j is only used to calibrate the threshold, default dependence is generated by the dependence structure of $A_1, \ldots, A_M$. Given this very general framework, there are still many ways to model the dependence between credits, of which we mention a few. A very popular approach is so-called times-to-default (TtD) modelling. For each credit, the basic modelling object is the random default time τ_j. It is generated by transforming the latent variable A_j via the marginal distribution function F_j. The famous Gaussian copula variant of this model, introduced by Li (2000), is obtained by choosing $A_1, \ldots, A_M$ as multivariate normal and setting

$$\tau_j = F_j^{-1}\left(\Phi(A_j)\right).$$

This specification immediately fits into the framework of equation (6) with $K_j = \Phi^{-1}(p_j)$. The TtD approach with a Gaussian copula is very tractable, but it also has fairly severe flaws. The dependence structure between credits is highly time-inhomogeneous, and the model does not produce realistic spread dynamics. It needs to be adjusted to match observed CDO tranche prices, i.e. fit to the “correlation smile”. Nevertheless, it has become a de facto market standard. A more dynamic approach is achieved by defining a credit index process X_j for each credit j: default occurs at the first time X_j hits a time dependent barrier. This was originally proposed by Black and Cox (1976) and has more recently been revived by Hull and White (2001). A somewhat different approach uses the Cox process framework. Defaults are generated by jumps of point processes which are given as time changes of independent Poisson processes via dependent stochastic hazard rates. However, to produce realistic CDO tranche spreads, one needs to introduce jumps in the hazard rates or other feedback effects, because correlated diffusions do not generate sufficient levels of default dependence.

Despite the existence of a multitude of modelling approaches, several mathematical ideas have proven themselves very powerful in dealing with the high dimensionality inherent to credit portfolio analysis: conditional independence, asymptotic methods and semi-asymptotic sensitivity calculations.

2.1. Conditional Independence

Conditional independence is the idea that, conditional on some mixing variable η the credits in the portfolio are independent. The simplest example of this is the one-factor Gaussian copula, where each credit's $\mathcal{N}(0,1)$ distributed latent variable A_j is given by

$$A_j = \beta_j \, Z_{Mkt} + \sqrt{1-\beta_j^2}\, Z_j, \tag{7}$$

and the variables $Z_{Mkt}, Z_1, \ldots, Z_M$ are i.i.d $\mathcal{N}(0,1)$. Conditional on the market factor Z_{Mkt}, all the credits are independent. In general, the conditional loss distribution is binomial and the unconditional loss distribution is obtained by integrating over η. Particularly in the one-factor framework, this is a straightforward numerical integration. A lot of research effort has been put into finding methods of computing the conditional loss distributions efficiently. Popular techniques include Fourier transforms, recursion techniques, as well as saddlepoint and other analytical approximations. The recursion method is particularly intuitive: for simplicity, we assume that each credit generates the same loss at default, so that the loss distribution can be expressed in integer multiples of an underlying loss unit. Denote the conditional default probability of credit j by p_j. For each $n \in \{0, 1, \ldots, M\}$, $L^{(n)}$ is the conditional portfolio loss after n credits have been added to the portfolio, and $p_k^{(n)} = P\left[L^{(n)} = k\right]$. The start of the recursion is clear, because $p_0^{(0)} = 1$ and $p_k^{(0)} = 0$ for $k > 0$. The recursion step is

$$p_{k+1}^{(n)} = p_{k+1}^{(n-1)}(1-p_n) + p_k^{(n-1)} p_n, \tag{8}$$

with $p_0^{(n)} = p_0^{(n-1)}(1-p_n)$. Note that each credit makes a “default” and a “survival” contribution. When pricing a tranche, we can pick out the slice of the loss distribution we are interested in, depending on the strikes of the tranche. To compute sensitivities, we can use equation (8) to quickly unwind a step of the recursion. Because conditional default probabilities can be very small, numerical stability is key. Hence one needs to anchor the recursion either at a zero loss or the maximum loss.

2.2. The LHP Approximation and Semi-Asymptotic Extensions

Another important stochastic tool for the analysis of credit portfolios is the so-called Large Homogeneous Portfolio (LHP) approximation. This was introduced in a Gaussian context by Vasicek (1987), see also Vasicek (2002). In the conditional independence framework, one assumes the portfolio consists of equally-weighted homogeneous assets with default probability $p(\eta)$. As the number M of credits tends to infinity, the fraction of credits defaulting converges to $p(\eta)$ by the Law of Large Numbers. For simplicity, assume that recovery rates are zero, then the fractional loss of the portfolio also converges to $p(\eta)$. In the LHP approximation, we assume that the conditional loss is actually $p(\eta)$. We are replacing the conditional loss distribution with its conditional expectation, i.e. matching the first moment of the conditional loss distribution with a point mass. The great advantage of the LHP approximation is that it often gives analytical formulae and is therefore very useful in developing intuition for a given dependence structure. Since the original work Vasicek (1987), Schönbucher (2004) has applied the LHP approximation to some of the Archimedean copula family, Schloegl and O'Kane (2005) have analyzed the Student-t copula. In the Gaussian case, we start from equation (7). The default probability p and the correlation parameter β are common across all credits, the default threshold C is given by $C = \Phi^{-1}(p)$. Assuming $\beta > 0$, we can compute the unconditional loss distribution to be

$$P\left[L \leq \theta\right] = \Phi\left(\frac{\sqrt{1-\beta^2}\,\Phi^{-1}(\theta) - \Phi^{-1}(p)}{\beta}\right).$$

Because the LHP approximation assumes that the portfolio is homogeneous, it is not well suited to computing sensitivities to changes in individual issuer characteristics. A way to deal with this is via so-called semi-asymptotic methods. In a given portfolio, we model the credit we are interested in exactly, while treating the rest of the portfolio asymptotically. Emmer and Tasche (2003) use this approach to compute risk capital contributions. Lehman Brothers has utilized this method for computing sensitivities in a model we call LHP plus one asset, or LH+ for short, cf. Greenberg et al. (2004). To compute a stop-loss transform $E\left[[L-K]^+\right]$ for a given strike K, it is possible to identify thresholds $A < B$ for the market factor which determine whether the single credit is relevant for crossing the strike or not. Extending the LHP analysis then gives a formula in terms of bi- and trivariate normal distributions.

$$\begin{aligned} E\left[[L-K]^+\right] &= K\Phi_{2,\beta_0}(C_0, A) + (N_0 - K)\,\Phi_{2,\beta_0}(C_0, B) \\ &\quad + N\left[\Phi_{2,\beta}(C, A) + \Phi_{3,\Sigma}(C_0, C, B) - \Phi_{3,\Sigma}(C_0, C, A)\right]. \end{aligned} \tag{9}$$

Computing the spread delta effectively entails differentiating equation (9) with respect to the individual credit's default threshold C_0. This reduces the trivariate normal distribution to a bivariate one, giving a very tractable formula for the spread sensitivity. Further computational enhancements can be achieved by modelling the conditional loss distribution more exactly. One would immediately think of a Central Limit Theorem argument, i.e. matching the first two moments of the conditional loss distribution by fitting a normal distribution, as has been proposed by Finger (1999). However, the Central Limit Theorem can easily be a false friend in credit modelling, as we are dealing with rare events and are often concerned with the tail of the distribution. In fact, we have found that fitting a simple two-point distribution to the conditional loss by matching moments gives better results than a Gaussian approximation in the LH+ context.

2.3. CDO^2 Pricing

An interesting application of the modelling techniques we have discussed is the pricing of CDO^2 structures. The fundamental underlying to such a structure is a large pool of credits, the main constraint here is number of entities liquidly traded in the CDS market. The individual credits are assigned to different miniportfolios. The miniportfolios do overlap and the weighting of a particular credit is specific to each miniportfolio. A bespoke CDO tranche is chosen for each miniportfolio, and these minitranches form the reference set of a new structure, the synthetic CDO^2. Finally, a bespoke supertranche is selected from the squared structure. Of course, losses to the supertranche depend on losses affecting the minitranches, which in turn depend on the joint default behaviour of the individual credits. The seller of supertranche protection covers the losses affecting the supertranche, just as in a standard CDO tranche. Similarly, the contractual spread paid to the seller is based on the outstanding notional of the supertranche. The loss L_i^{tr} to the ith minitranche is a function of the loss to the ith miniportfolio as shown in equation (4). The percentage loss L^{sp} to the superportfolio is given as

$$L^{sp} = \frac{\sum_{i=1}^{N} N^{(i)} L_i^{tr}}{N_{tot}}.$$

Finally, the supertranche loss is given by

$$L^{st} = \frac{[L^{sp} - K_{st}]^+ - [L^{sp} - (K_{st} + w_{st})]^+}{w_{st}}. \tag{10}$$

Equation (10) shows that the supertranche loss is a compound option on the joint distribution of all the miniportfolio losses $L_1, \ldots, L_N$. This poses two challenges. Even if the individual credits are independent, $L_1, \ldots, L_N$ are not because of the overlap of the miniportfolios. Also, even if one has a very tractable model for the joint distribution of $L_1, \ldots, L_N$, one still needs to price a compound option. Naively treating the minitranches as effective CDS with some correlation is doomed to failure, as it is unclear which correlation to use. This is because tranches are highly non-linear payouts and some of the dependence stems from contagion between credits (this is captured by a factor model), whereas another part stems from the overlap between miniportfolios. Finally, the effective correlation is a function of the minitranche subordination. Senior tranches are more exposed to systemic risk, hence behave in a more correlated fashion.

One method to price CDO^2 tranches is to extend the recursion technique to higher dimensions, as shown in Baheti et al. (2005). Let us assume for simplicity that we are in the two-dimensional case, i.e. the superportfolio contains two minitranches. The conditional probability of joint losses after n credits is denoted by $p^n_{k_1,k_2}$. We add credits recursively with default probability π_j and loss weights of λ_i^j. Each credit makes a survival and a default contribution. We have both if $k_1 \geq \lambda_1^{n+1}$ and $k_2 \geq \lambda_2^{n+1}$. In this case the recursion is

$$p^{n+1}_{k_1,k_2} = (1 - \pi_{n+1})\, p^n_{k_1,k_2} + \pi_{n+1}\, p^n_{k_1-\lambda_1^{n+1},k_2-\lambda_2^{n+1}}.$$

Other points of the joint loss distribution only have a survival contribution

$$p^{n+1}_{k_1,k_2} = (1 - \pi_{n+1})\, p^n_{k_1,k_2}.$$

The main limitation of this approach is the fact that the probability space, and hence memory requirements, grow exponentially with the number of miniportfolios. The method is useful for relatively medium scale structures (around 7 minitranches), and also for many different similar products. For dealing with higher dimensional cases, other methods are needed. One approach is to condition on the market factor, simulate the joint loss distribution of $L_1, \ldots, L_N$, and evaluate the compound option via Monte Carlo. The moments of the conditional distribution are easy to calculate, because the individual credits are independent.

$$E\left[L_i \middle| Z\right] = \sum_{k=1}^{M} (1 - R_k) w_{i,k}\, p_k(Z) \tag{11}$$

$$\text{cov}\left(L_i, L_j \middle| Z\right) = \sum_{k=1}^{M} (1 - R_k)^2 w_{i,k}\, w_{j,k}\, p_k(Z)\left(1 - p_k(Z)\right). \tag{12}$$

One way of accelerating the conditional Monte Carlo simulation is to use equation (11) and (12) to fit a multivariate normal distribution to the conditional losses. However, as mentioned before, one has to be wary of the Central Limit Theorem as a false friend in this type of application.

3. A BRIEF OUTLOOK

In the previous section we have detailed some of the stochastic techniques currently useful in credit derivatives pricing and have illustrated some of their applications. We have not discussed the very important topic of the correlation smile, i.e. the deviation between the market pricing of tranches and the simple Gaussian copula. Practitioners are very much searching for models which fit the observed market prices in the best possible manner. Also, in terms of modelling, one really needs to move beyond the simple times-to-default framework. It is too static, induces counterintuitive time-inhomogeneities and hence produces unrealistic spread dynamics. In fact, the challenge to all researchers in the credit derivatives field is to think more dynamically about modelling credit risk and portfolio credit derivatives. This will allow us to develop models which produce more realistic spread dynamics.

References

P. Baheti, R. Mashal, M. Naldi, and L. Schloegl. Squaring factor copula models. *Risk*, to appear, 2005.

F. Black and J.C. Cox. Valuing corporate securities: Some effects of bond indenture provisions. *Journal of Finance*, 31:351–367, 1976.

S. Emmer and D. Tasche. Calculating credit risk capital charges with the one-factor model. working paper, 2003.

C.C. Finger. Conditional approaches for CreditMetrics portfolio distributions. *CreditMetrics Monitor*, pages 14 – 33, April 1999.

A. Greenberg, D. O'Kane, and L. Schloegl. LH+: A fast analytical model for CDO hedging and risk management. *Lehman Brothers Quantitative Credit Research Quarterly*, 2004-Q2, 2004.

J.C. Hull and A. White. Valuing credit default swaps II: Modeling default correlations. *Journal of Derivatives*, 8(3):12–22, Spring 2001.

D.X. Li. On default correlation: a copula function approach. *Journal of Fixed Income*, 9:43–54, 2000.

L. Schloegl and D. O'Kane. A note on the large homogeneous portfolio approximation with the Student-t copula. *Finance and Stochastics*, to appear, 2005.

P.J. Schönbucher. Taken to the limit: Simple and not-so-simple loan loss distributions. *Wilmott magazine*, 2004.

O. Vasicek. Probability of loss on loan portfolio. working paper, KMV Corporation, 1987.

O. Vasicek. Loan portfolio value. *Risk*, 15(12):160–162, December 2002.